What is claimed is:

1. In a file server system having a clock for producing a clock time and a processor for servicing client requests for access to a file, the processor having a timer for measuring a time interval, a method comprising:

the processor obtaining the clock time from the clock, and beginning measurement of the time interval with the timer, and

the processor responding to a request from a client for an asynchronous write to the file by performing an asynchronous write operation with respect to the file, and determining a file-modification time that is a function of the clock time having been obtained from the clock and the time interval measured by the timer, the file-modification time indicating a time of modification of the file by the asynchronous write operation.

2. The method as claimed in claim 1, wherein the file-modification time is a sum of the clock time having been obtained from the clock and the time interval measured by the timer.

3. The method as claimed in claim 1, which includes the processor acknowledging the request from the client for an asynchronous write to the file by returning to the client the file-modification time.

4. The method as claimed in claim 1, which further includes the processor receiving an updated value for the file-modification time after the processor has determined a value

for the file-modification time, the processor comparing the updated value to the value that the processor has determined for the file-modification time, and upon finding that the updated value for the file-modification time is greater than the value that the processor has determined for the file-modification time, then the processor resetting the timer and using the updated value for the file-modification time in lieu of the clock time obtained from the clock.

5. The method as claimed in claim 4, wherein the processor stores the clock time having been obtained from the clock in a memory location local to the processor, and the processor uses the updated value for the file-modification time in lieu of the clock time obtained from the clock by replacing the clock time having been obtained from the clock and stored in the memory location local to the processor with the updated value for the file-modification time.

6. The method as claimed in claim 1, which further includes the processor receiving an updated value for the file-modification time after the processor has determined a value for the file-modification time, the processor comparing the updated value for the file-modification time to the value that the processor has determined for the file-modification time, and upon finding the updated value for the file-modification time is less than the value that the processor has determined for the file-modification time, then the processor ignoring the updated value for the file-modification time.

7. In a file server system having a first processor and a second processor for servicing client requests for access to a file, the first processor having a clock producing a clock time, and the second processor having a timer for measuring a time interval, a method comprising:

the second processor responding to a first request from a client for an asynchronous write to the file by obtaining the clock time from the clock of the first processor, beginning measurement of the time interval with the timer, performing a first asynchronous write operation with respect to the file, and using the clock time obtained from the clock of the first processor as a first file-modification time, the first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and thereafter

the secondary processor responding to a second request from the client for an asynchronous write to the file by performing a second asynchronous write operation with respect to the file, and determining a second file-modification time that is a function of the clock time obtained from the clock of the first processor and the time interval measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

8. The method as claimed in claim 7, wherein the file-modification time is a sum of the clock time having been obtained from the clock and the time interval measured by the timer.

9. The method as claimed in claim 7, which includes:

the second processor acknowledging the first request from the client for an asynchronous write to the file by returning to the client the first file-modified time for the file; and

the second processor acknowledging the second request from the client for an asynchronous write to the file by returning to the client the second file-modified time for the file.

10. The method as claimed in claim 7, which includes the second processor responding to a request from the client to commit results of the second asynchronous write operation by sending the second file-modification time to the first processor.

11. In a file server system having a first processor and a second processor for servicing client requests for access to a file, the first processor having a clock producing a clock time, and the second processor having a timer for measuring a time interval, a method comprising:

the second processor responding to a first request from a client for an asynchronous write to the file by obtaining the clock time from the clock of the first processor, beginning measurement of the time interval with the timer, performing a first asynchronous write operation with respect to the file, and using the clock time obtained from the clock of the first processor as a first file-modification time, the first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and thereafter

the second processor receiving from the first processor an updated value for the file-modification time, the second processor comparing the updated value for the file-modification time to the first file-modification time, and upon finding that the updated value is greater than the first file-modification time, the second processor resetting the timer; and thereafter

the second processor responding to a second request from the client for an asynchronous write to the file by performing a second asynchronous write operation with respect to the file, and determining a second file-modification time that is a sum of the updated value for the file-modification time and the time interval measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

12. In a file server system having a primary processor managing metadata of a file, and a secondary processor responding to requests from a client for access to the file, the primary processor having a clock producing a clock time, and the secondary processor having a timer for measuring a time interval, a method comprising:

the secondary processor responding to a first asynchronous write request from the client for writing to the file by obtaining attributes of the file and the clock time from the primary processor, storing the attributes of the file in a cache local to the secondary processor and using the file attributes to perform a first asynchronous write operation with respect to the file, and beginning measurement of the time interval with the timer, and thereafter

the secondary processor responding to a second asynchronous write request from the client for writing to the file by using the attributes of the file in the cache local to the secondary processor to perform a second asynchronous write operation with respect to the file, and determining a file-modification time that is a function of the clock time having been obtained from the clock of the primary processor and the interval measured by the timer, the file-modification time indicating a time of modification of the file by the second asynchronous write operation.

13. The method as claimed in claim 12, wherein the file-modification time is a sum of the clock time having been obtained from the clock and the time interval measured by the timer.

14. The method as claimed in claim 12, which includes:

the secondary processor acknowledging the second asynchronous write request from the client by returning to the client the file-modification time as the time when the file was modified by the second asynchronous write operation.

15. The method as claimed in claim 12, which includes:

the secondary processor responding to a request from the client to commit results of the second asynchronous write operation by sending a flush request to the primary processor, the flush request including the file-modification time.

16. The method as claimed in claim 15, which includes the primary processor sending the file-modification time to another client caching attributes for the file.

17. The method as claimed in claim 12, which includes the secondary processor receiving from the primary processor an updated value for the file-modification time after the secondary processor has completed the second asynchronous write operation, the secondary processor comparing the updated value for the file-modification time to the last value for the file-modification time determined by the secondary processor, and upon finding that the updated value for the file-modification time is greater than the last value for the file-modification time determined by the secondary processor, the secondary processor resetting the timer, and using the updated value for the file-modification time in lieu of the clock time having been obtained from the primary processor, and using the updated value for the file-modification time as the most recent value of the file-modification time.

18. The method as claimed in claim 12, which includes the secondary processor receiving from the primary processor an updated value for the file-modification time after the secondary processor has completed the second asynchronous write operation, the secondary processor comparing the updated value for the file-modification time to the last value for the file-modification time determined by the secondary processor, and upon finding that the updated value for the file-modification time is less than the last value for the file-modification time determined by the secondary processor, the secondary processor ignoring the updated value for the file-modification time.

19. In a network file server having a plurality of data mover computers for servicing client requests for access to a file, and a cached disk array for storing data of the file, the data mover computers being coupled to the cache disk array for accessing the data of the file, the data mover computers including a primary data mover computer managing metadata of the file, and a secondary data mover computer that requests metadata of the file from the primary data mover computer, the primary data mover computer having a clock producing a clock time, and the secondary data mover computer having a timer for measuring a time interval, a method comprising:

the secondary data mover computer responding to a first asynchronous write request from a client for writing to the file by obtaining attributes of the file and the clock time from the primary data mover computer, storing the attributes of the file in a cache local to the secondary data mover computer and using the file attributes to perform a first asynchronous write operation with respect to the file, and using the clock time as a first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and thereafter

the secondary data mover computer responding to a second asynchronous write request from the client for writing to the file by using the attributes of the file in the cache local to the secondary data mover computer to perform a second asynchronous write operation with respect to the file, and determining a second file-modification time that is a function of the clock time having been obtained from the primary data mover and the time interval measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

20. The method as claimed in claim 19, wherein the second file-modification time is a sum of the clock time having been obtained from the primary data mover and the time interval measured by the timer.

21. The method as claimed in claim 19, wherein:

the secondary data mover computer uses the clock time as a first file-modification time by acknowledging the first asynchronous write request from the client by returning to the client the clock time as the time when the file was modified by the first asynchronous write operation, and

the secondary data mover computer acknowledges the second asynchronous write request from the client by returning to the client the second file-modification time as the time when the file was modified by the second asynchronous write operation.

22. The method as claimed in claim 19, which includes:

the secondary data mover computer responding to a request from the client to commit results of the second asynchronous write operation by sending a flush request to the primary data mover computer, the flush request including the second file-modification time.

23. The method as claimed in claim 22, which includes the primary data mover computer sending the second file-modification time to another client caching attributes for the file.

24. The method as claimed in claim 19, which includes the secondary data mover computer receiving from the primary data mover computer an updated value for the file-modification time for the file after the secondary data mover computer has completed the first asynchronous write operation, the secondary data mover computer comparing the updated value for the file-modification time for the file to the last value determined by the secondary data mover for the file-modified time for the file, and upon finding that the updated value for the file-modification time for the file is greater than the last value determined by the secondary data mover for the file-modified time for the file, the secondary data mover computer resetting the timer, using the updated value for the file-modification time in lieu of the clock time having been obtained from the primary data mover computer, and using the updated value for the file-modification time for the file as the most recent value for the file-modification time for the file.

25. The method as claimed in claim 19, which includes the secondary data mover computer receiving from the primary data mover computer an updated value for the file-modification time for the file after the secondary data mover computer has completed the first asynchronous write operation, the secondary data mover computer comparing the updated value for the file-modification time for the file to the last value determined by the secondary data mover for the file-modified time for the file, and upon finding that the updated value for the file-modification time for the file is less than the last value determined by the secondary data mover for the file-modified time for the file, the secondary data mover computer ignoring the updated value for the file-modification time.

26. A file server system having a clock for producing a clock time and a processor for servicing client requests for access to a file, the processor having a timer for measuring a time interval;

the processor being programmed for obtaining the clock time from the clock, and beginning measurement of the time interval with the timer, and

the processor being programmed for responding to a request from a client for an asynchronous write to the file by performing an asynchronous write operation with respect to the file, and determining a file-modification time that is a function of the clock time having been obtained from the clock and the time interval measured by the timer, the file-modification time indicating a time of modification of the file by the asynchronous write operation.

27. The file server system as claimed in claim 26, wherein the file-modification time is a sum of the clock time having been obtained from the clock and the time interval measured by the timer.

28. The file server system as claimed in claim 26, wherein the processor is programmed to acknowledging the request from the client for an asynchronous write to the file by returning to the client the file-modification time.

29. The file server system as claimed in claim 26, wherein the processor is programmed for receiving an updated value for the file-modification time after the

processor has determined a value for the file-modification time, comparing the updated value for the file-modification time to the value that the processor has determined for the file-modification time, and upon finding the updated value for the file-modification time is greater than the value that the processor has determined for the file-modification time, resetting the timer and using the updated value for the file-modification time in lieu of the clock time having been obtained from the clock.

30. The file server system as claimed in claim 29, wherein the processor is programmed for storing the clock time having been obtained from the clock in a memory location local to the processor, and for using the updated value for the file-modification time in lieu of the clock time having been obtained from the clock by replacing the clock time stored in the memory local to the processor with the updated value for the file-modification time.

31. The file server system as claimed in claim 26, wherein the processor is programmed for receiving an updated value for the file-modification time after the secondary processor has determined a value for the file-modification time, comparing the updated value to the value that the processor has determined for the file-modification time, and ignoring the updated value for the file-modification time upon finding the updated value for the file-modification time is less than the value that the secondary processor has determined for the file-modification time.

32. A file server system comprising:

a first processor and a second processor for servicing client requests for access to a file, the first processor having a clock for producing a clock time, and the second processor having a timer for measuring a time interval;

the second processor being programmed for responding to a first request from a client for an asynchronous write to the file by obtaining the clock time from the clock of the first processor, beginning measurement of the time interval with the timer, performing a first asynchronous write operation with respect to the file, and using the clock time obtained from the clock of the first processor as a first file-modification time, the first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and the second processor being programmed for responding to a second request from the client for an asynchronous write to the file by performing a second asynchronous write operation with respect to the file, and determining a second file-modification time that is a function of the clock time obtained from the clock of the first processor and the time interval measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

33. The file server system as claimed in claim 32, wherein the second file-modification time is a sum of the clock time obtained from the clock and the time interval measured by the timer.

34. The file server system as claimed in claim 32, wherein:

the second processor is programmed to use the clock time obtained from the clock of the first processor as a first file-modification time by acknowledging the first request from the client for an asynchronous write to the file by returning to the client the clock time obtained from the clock of the first processor as the time when the file was modified by the first asynchronous write operation, and

the second processor is programmed to acknowledge the second request from the client for an asynchronous write to the file by returning to the client the second file-modification time as the time when the file was modified by the second asynchronous write operation.

35. The file server system as claimed in claim 32, wherein the second processor is programmed for responding to a request from the client to commit results of the second asynchronous write operation by sending the second file-modification time to the first processor.

36. A file server system comprising:

a first processor and a second processor for servicing client requests for access to a file, the first processor having a clock for producing a clock time, and the second processor having a timer for measuring a time interval;

the second processor being programmed for responding to a first request from a client for an asynchronous write to the file by obtaining the clock time from the clock of the first processor, beginning measurement of the time interval with the timer, and

performing a first asynchronous write operation with respect to the file, and using the clock time obtained from the clock of the first processor as a first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and

the second processor being programmed for receiving from the first processor an updated value for the file-modification time, for comparing the updated value to the first file-modification time, and upon finding that the updated value is greater than the first file-modification time, for resetting the timer; and

the second processor being programmed to respond to a second request from the client for an asynchronous write to the file by performing a second asynchronous write operation with respect to the file, and determining a second file-modification time that is a sum of the updated value for the file-modification time and the time measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

37. A file server system comprising:

a primary processor managing metadata of a file, and a secondary processor responding to requests from a client for access to the file, the primary processor having a clock for producing a clock time, and the secondary processor having a timer for measuring a time interval;

the secondary processor being programmed for responding to a first asynchronous write request from the client for writing to the file by obtaining attributes of the file and the clock time from the primary processor, storing the attributes of the file in a cache local to the secondary processor and using the file attributes to perform a first

asynchronous write operation with respect to the file, and beginning measurement of the time interval with the timer; and

the secondary processor being programmed for responding to a second asynchronous write request from the client for writing to the file by using the attributes of the file in the cache local to the secondary processor to perform a second asynchronous write operation with respect to the file, and determining a file-modification time that is a function of the clock time having been obtained from the clock of the primary processor and the time interval measured by the timer, the file-modification time indicating a time of modification of the file by the second asynchronous write operation.

38. The file server system as claimed in claim 37, wherein the file-modification time is a sum of the clock time having been obtained from the primary processor and the time interval measured by the timer.

39. The file server system as claimed in claim 37, wherein:

the secondary processor is programmed for acknowledging the second asynchronous write request from the client by returning to the client the file-modification time as the time when the file was modified by the second asynchronous write operation.

40. The file server system as claimed in claim 37, wherein the secondary processor is programmed for responding to a request from the client to commit results of the second asynchronous write operation by sending a flush request to the primary processor, the flush request including the file-modification time.

41. The file server system as claimed in claim 40, wherein the primary processor is programmed to send the file-modification time to other clients caching attributes for the file.

42. The file server system as claimed in claim 37, wherein the secondary processor is programmed for receiving from the primary processor an updated value for the file-modification time after the secondary processor has completed the second asynchronous write operation, for comparing the updated value for the file-modification time to the last value for the file-modification time determined by the secondary processor, and upon finding that the updated value for the file-modification time is greater than the last value for the file-modification time determined by the secondary processor, for resetting the timer, and using the updated value for the file-modification time in lieu of the clock time having been obtained from the primary processor, and using the updated value for the file-modification time as the most recent value of the file-modification time.

43. The file server system as claimed in claim 37, wherein the secondary processor is programmed for receiving from the primary processor an updated value for the file-modification time after the secondary processor has completed the second asynchronous write operation, for comparing the updated value for the file-modification time to the last value for the file-modification time determined by the secondary processor, and upon finding that the updated value for the file-modification time is less than the last value for

the file-modification time determined by the secondary processor, for ignoring the updated value for the file-modification time.

44. A network file server comprising:

a plurality of data mover computers for servicing client requests for access to a file, and a cached disk array for storing data of the file, the data mover computers being coupled to the cached disk array for accessing the data of the file, the data mover computers including a primary data mover computer programmed for managing metadata of the file, and a secondary data mover computer programmed for requesting metadata of the file from the primary data mover computer, the primary data mover computer having a clock for producing a clock time, and the secondary data mover computer having a timer for measuring a time interval;

the secondary data mover computer being programmed for responding to a first asynchronous write request from a client for writing to the file by obtaining attributes of the file and the clock time from the primary data mover computer, storing the attributes of the file in a cache local to the secondary data mover computer and using the file attributes to perform a first asynchronous write operation with respect to the file, beginning measurement of the time interval with the timer, and using the clock time as a first file-modification time, the first file-modification time indicating a time of modification of the file by the first asynchronous write operation; and

the secondary data mover computer being programmed for responding to a second asynchronous write request from the client for writing to the file by using the attributes of the file in the cache local to the secondary data mover computer to perform a second

asynchronous write operation with respect to the file, and determining a second file-modification time that is a function of the clock time having been obtained from the primary data mover and the time interval measured by the timer, the second file-modification time indicating a time of modification of the file by the second asynchronous write operation.

45. The network file server as claimed in claim 44, wherein the second file-modification time is a sum of the clock time having been obtained from the primary data mover and the time interval measured by the timer.

46. The network file server as claimed in claim 44, wherein:

the secondary data mover computer is programmed for using the clock time as a first file-modification time by acknowledging the first asynchronous write request from the client by returning to the client the clock time as the time when the file was modified by the first asynchronous write operation, and

the secondary data mover computer is programmed for acknowledging the second asynchronous write request from the client by returning to the client the second file-modification time as the time when the file was modified by the second asynchronous write operation.

47. The network file server as claimed in claim 44, wherein the secondary data mover computer is programmed for responding to a request from the client to commit results of

the second asynchronous write operation by sending a flush request to the primary data mover computer, the flush request including the second file-modification time.

48. The network file server as claimed in claim 47, wherein the primary data mover computer is programmed for sending the second file-modification time to other clients caching attributes for the file.

49. The network file server as claimed in claim 44, wherein the secondary data mover computer is programmed for receiving from the primary data mover computer an updated value for the file-modification time for the file after the secondary data mover computer has completed the first asynchronous write operation, for comparing the updated value for the file-modification time for the file to last value determined by the secondary data mover for the file-modified time for the file, and upon finding that the updated value for the file-modification time for the file is greater than the last value determined by the secondary data mover for the file-modified time for the file, for resetting the timer, using the updated value for the file-modification time for the file in lieu of the clock time having been obtained from the primary data mover computer, and using the updated value for the file-modified time for the file as the most recent value for the file-modification time for the file.

50. The network file server as claimed in claim 44, wherein the secondary data mover computer is programmed for receiving from the primary data mover computer an updated value for the file-modification time for the file after the secondary data mover computer

has completed the first asynchronous write operation, for comparing the updated value for the file-modification time for the file to last value determined by the secondary data mover for the file-modified time for the file, and upon finding that the updated value for the file-modification time for the file is less than the last value determined by the secondary data mover for the file-modified time for the file, for ignoring the updated value for the file-modification time for the file.